Inverse Reinforcement Learning in Swarm Systems

Adrian Šošić and **Wasiur R. KhudaBukhsh** and **Abdelhak M. Zoubir** and **Heinz Koeppl**

arXiv:1602.05450v1 [stat.ML] 17 Feb 2016

Abstract. Inverse reinforcement learning (IRL) is the problem of recovering a system's latent reward function from observed system behavior. In this paper, we concentrate on IRL in homogeneous large-scale systems, which we refer to as swarms. We show that, by exploiting the inherent homogeneity of a swarm, the IRL objective can be reduced to an equivalent single-agent formulation of constant complexity, which allows us to decompose a global system objective into local subgoals at the agent-level. Based on this finding, we reformulate the corresponding optimal control problem as a fix-point problem pointing towards a symmetric Nash equilibrium, which we solve using a novel heterogeneous learning scheme particularly tailored to the swarm setting. Results on the Vicsek model and the Ising model demonstrate that the proposed framework is able to produce meaningful reward models from which we can learn near-optimal local controllers that replicate the observed system dynamics.

1 Introduction

Emergence and the ability of self-organization are fascinating characteristics of natural systems with interacting agents. Without a central controller, these systems are inherently robust to failure while, at the same time, they show remarkable large-scale dynamics that allow for fast adaptation to changing environments. Interestingly, for large system sizes, it is usually not the complexity of the individual agent but their (local) coupling that predominantly gears the final system dynamics. It has been shown [23, 31], in fact, that even relatively simple local dynamics can result in various kinds of higher-order complexity at a global scale when coupled through a network with many agents. In this work, we aim for a deeper understanding of the interplay between these two scales by answering the question how a global system objective can be decomposed into local subgoals at the agent-level through the principle of inverse reinforcement learning (IRL). In particular, we show how the homogeneity of a system can be exploited to reduce the complexity of the resulting IRL formulation.

While there has been a significant increase of scientific research in the field of IRL in the last decade, only few attempts have been made to transfer the concept of IRL to problems with more than one agent or, in particular, large-scale systems. The authors of [9], for instance, presented a multi-agent IRL-like framework based on mechanism design. Unfortunately, the proposed framework can only be used to refine a given reward model in order to promote a certain system behavior, but not to learn the necessary reward structure from demonstration data. In [20], the original IRL formulation (see [22]) has been extended to non-cooperative multi-agent problems to learn a joint reward model that can explain the system behavior at a global scale. While the presented approach does not benefit from a homogeneous agent architecture, the authors furthermore assume all agents in the network to be controlled by a central mediator, an assumption which is clearly inappropriate for self-organizing systems. A decentralized solution was later presented in [25] but, again, the proposed algorithm has been designed for general non-cooperative scenarios and is hence unable to exploit potential symmetries of the system.

In contrast to previous work on multi-agent IRL, we do not consider the general class of multi-agent systems but deliberately restrict ourselves to the rich sub-class of homogeneous systems. More specifically, we focus on large-scale networks, arguing that these systems deserve a separate treatment due to their specific structure and properties. The rationale is that, unlike in most single-agent or small-scale problems involving complicated agent architectures, the complexity of such networks typically arises from the sheer amount of local interactions, making the resulting learning problems particularly hard to solve [2, 7, 27]. However, finding a way to efficiently handle these systems could enable a number of promising future technologies with numerous potential applications (e.g. nanomedicine [10], the study of collective animal behavior [6, 8, 31], or the design of complex distributed systems like programmable matter [11], distributed sensor networks [15, 32] or self-assembly structures [34]).

In the following, our contribution is twofold: first, we provide an efficient IRL solution for the homogeneous multi-agent setting whose complexity stays constant for increasing system sizes. In order to arrive at such a solution, we exploit the homogeneity property of the systems and show that we can reduce the IRL formulation to an equivalent single-agent problem. Second, we investigate how the learned reward function can be used in the context of distributed system design. In particular, we address the question how a given behavior, perhaps planned on the global level, can be reconstructed from the agent level when there are only local observations available.

The outline of the paper is as follows: Section 2 shortly recapitulates the basic idea of IRL in the single-agent case. In Section 3, we introduce our model and discuss the IRL objective in the context of homogeneous systems, focusing on the task of value estimation. Section 4 addresses the problem of optimal control and presents a new model-free learning strategy for homogeneous systems. In Section 5, we present empirical results of our framework on the Vicsek model and the Ising model. Finally, we conclude the paper in Section 6 and discuss some of the open problems.

2 Inverse Reinforcement Learning

Inverse reinforcement learning refers to the problem of recovering a system's latent objective from observed system behavior, typically provided in the form of trajectory data. Under the Markov decision processes (MDP) formalism [28], this is equivalent to finding a reward function that is able to explain the given observations. The idea of IRL is particularly appealing in the context of reinforcement learning itself, as it offers a more principled approach to learning optimal controllers by inferring a suitable reward model from ex-

pert demonstrations as compared to using hand-crafted reward functions [12, 21]. However, IRL can be also applied for system identification as a powerful alternative to other inference techniques since the learned reward function gives direct insights into the working principle of a system and, as it captures the system objective compactly, it can serve as a basis for further system analysis (such as preference elicitation [26]).

The idea of IRL was first presented in [22] and several other approaches were proposed thereafter (e.g. [36, 17]; see [35] for an overview). However, IRL has not been well-explored for multi-agent systems which is why, in this paper, we attempt to build a sound mathematical framework for IRL in the homogeneous setting. In particular, we exploit the fact that many existing IRL frameworks are value-based[1] and share a common generic algorithmic form [19] (see below) which invokes three main steps: 1) policy update 2) value estimation and 3) reward update. The third step is generally algorithm-specific and hence not of particular interest for this paper. In fact, we show that, in a homogeneous setup, both the policy update and the value estimation task can be essentially reduced to equivalent single-agent problems, meaning that the same basic algorithmic form applies and, therefore, the final multi-agent algorithm can be instantiated using any existing value-based single-agent method. For the remainder of the paper, we thus focus on the first two steps and explain how they can be efficiently realized in a homogeneous system.

Algorithm 1: GENERIC IRL

Input: MDP without reward function R, expert demonstrations $\mathcal{D}$

1: Initialize reward function $R^{\{0\}}$
for $i = 0, 1, 2, ...$
2: **Policy update:** Find an optimal policy $\pi^{\{i\}}$ given $R^{\{i\}}$
3: **Value estimation:** Compute $V^{\{i\}}$ given $\pi^{\{i\}}$ and $R^{\{i\}}$
4: **Reward update:** Based on $V^{\{i\}}$ and $\mathcal{D}$, compute $R^{\{i+1\}}$
end for

3 The swarMDP Model

While it is straightforward to generalize the principles of both RL and IRL to multi-agent systems by simply defining appropriate joint state and action spaces, respectively, such an extension would suffer from an exponential growth of state and action configurations, rendering most resulting models computationally intractable. However, in many practical settings, we can benefit from structure or symmetries that are inherent in the system in order to reduce the complexity of the final model. This is especially true for most large-scale systems that can be encountered in nature (e.g. herds of animals), where all individuals belong to the same species so that they can be (at least approximately) treated as identical copies of a single generic agent prototype.

As an analogy to their biological counterparts, we will in the following refer to systems with this property as *swarms*. Formally, we can construct a swarm system from a set of locally coupled MDPs with identical model specifications. An agent in such a swarm, which we call a *swarming agent*, is defined as a tuple $\mathbb{A} := (\mathcal{S}, \mathcal{O}, \mathcal{A}, T, R, \pi)$, where:

- $\mathcal{S}, \mathcal{O}, \mathcal{A}$ are finite sets of local states, observations and actions.
- $T : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$ is the local transition model of the agent. More specifically, $T(z, a, \tilde{z})$ denotes the conditional probability that the agent reaches state $\tilde{z}$ after performing action a at state z.
- $R : \mathcal{O} \to \mathbb{R}$ is an agent-level reward function.
- $\pi : \mathcal{O} \to \mathcal{A}$ is the local policy of the agent which later serves as the decentralized control law of the swarm.

Please note that we only consider the class of reactive policies, i.e. π is a function of the current observation only and not of the entire history of observations. In fact, the swarm systems considered in this paper (Section 5) are purely reactive. However, an extension to more general policy representations is straightforward but involves a more complicated policy update step (see e.g. [18]).

With the definition of the swarming agent at hand, we may now define a swarm as a tuple $(N, \mathbb{A}, \xi)$, where

- N is the number of agents in the swarm.
- $\mathbb{A}$ is a swarming agent prototype as defined above.
- $\xi : \mathcal{S}^N \to \mathcal{O}^N$ is the (global) observation model of the system defining which parts of the swarm state $s \in \mathcal{S}^N$ can be observed by whom. More precisely, $\xi(s) = (\xi^{(1)}(s), \dots, \xi^{(N)}(s))$ is the ordered set of local observations passed on to the agents at state s.

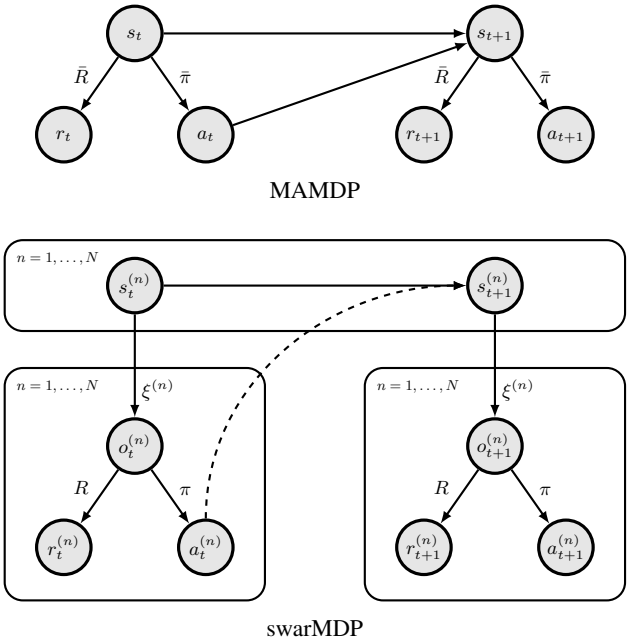

Figure 1: MAMDP and swarMDP as graphical models, the latter visualized using plate notation. As opposed to the more general MAMDP, a swarMDP explicitly defines a local policy $\pi : \mathcal{O} \to \mathcal{A}$ and a local reward function $R : \mathcal{O} \to \mathbb{R}$, both operating at the agent-level through the observation model ξ. The dashed arrow indicates connections between nodes belonging to the same agent.

In particular, we assume that the agents have no access to their local states $s^{(n)} \in \mathcal{S}$ but only to their local observations $o^{(n)} = \xi^{(n)}(s) \in \mathcal{O}$. Further, the observation model ξ is assumed to ensure homogeneity of the agents in the sense that agent n would receive agent m's observation (and vice versa) if they interchanged their local states. Mathematically, this means that any permutation of $s \in \mathcal{S}^N$ must result in the same permutation of $\xi(s)$. For example, in a swarm of fish, the local observation could be simply the alignment of a fish to its

[1] This can be easily explained by the fact that the IRL objective is most naturally stated in terms of value functions (see [22]).

immediate neighbors (see also Section 5.1). The graphical description of the model, in the sequel referred to as a *swarMDP*, is shown in Figure 1. As compared to a general multi-agent Markov decision process (MAMDP), our model explicitly obeys two key properties:

Homogeneity: All agents are identical in the sense that they share the same dynamics, reward function, action space and observation abilities.

Locality: The agents can observe and interact only with parts of the system that are close in terms of some distance metric as implied by the observation model ξ. As an immediate consequence, their decisions depend on their current neighborhood only and not on the whole swarm state, which considerably reduces the size of the policy space. Accordingly, we assume that the agents receive their rewards solely based on their local observations, which is in line with the concept of an internal reward system or mental state [3].

Due to the locality property, the joint dynamics of the system locally decouple (see next section), resulting in a significantly reduced model complexity. Moreover, the homogeneity property allows us to arrive at a scalable solution to the IRL problem (Section 3.3) and to efficiently solve the corresponding control problem (Section 4).

3.1 Private Value & Bellman Optimality

Analogous to the single-agent case [28], we define the *private value* of an agent n at a swarm state $s \in \mathcal{S}^N$ under policy π as the expected sum of discounted rewards accumulated by that agent over time, given that all agents execute π,

$$V^{(n)}(s \mid \pi) \coloneqq \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k R(\xi^{(n)}(s_{t+k})) \mid \pi, s_t = s\right], \quad (1)$$

where expectation is taken with respect to the stochastic state sequence starting from s, and γ denotes a discount factor in $[0, 1)$. Please note that due to the time-homogeneity of the local transition model T this definition of value is independent of any particular starting time t. Denoting further by $Q^{(n)}(s, a \mid \pi)$ the state-action value of agent n at state s for the case that all agents execute policy π, except for agent n who performs action a once and follows π thereafter, the following *Bellman equations* are obtained:

$$V^{(n)}(s \mid \pi) = R(\xi^{(n)}(s)) + \gamma \sum_{\tilde{s} \in \mathcal{S}^N} P(\tilde{s} \mid s, \pi) V^{(n)}(\tilde{s} \mid \pi),$$

$$Q^{(n)}(s, a \mid \pi) = R(\xi^{(n)}(s)) + \gamma \sum_{\tilde{s} \in \mathcal{S}^N} P^{(n)}(\tilde{s} \mid s, a, \pi) V^{(n)}(\tilde{s} \mid \pi).$$

Here, $P(\tilde{s} \mid s, \pi)$ denotes the probability of reaching swarm state $\tilde{s}$ from s when *every* agent performs policy π and, analogously, by $P^{(n)}(\tilde{s} \mid s, a, \pi)$, we mean the probability of reaching swarm state $\tilde{s}$ from state s if agent n chooses action a and *all other* agents execute policy π. Explicitly, for $s = (s^{(1)}, s^{(2)}, \ldots, s^{(N)}), \tilde{s} = (\tilde{s}^{(1)}, \tilde{s}^{(2)}, \ldots, \tilde{s}^{(N)}) \in \mathcal{S}^N$,

$$P(\tilde{s} \mid s, \pi) = \prod_{n=1}^{N} T(s^{(n)}, \pi(o^{(n)}), \tilde{s}^{(n)}),$$

$$P^{(n)}(\tilde{s} \mid s, a, \pi) = T(s^{(n)}, a, \tilde{s}^{(n)}) \prod_{\substack{l=1 \\ l \neq n}}^{N} T(s^{(l)}, \pi(o^{(l)}), \tilde{s}^{(l)}).$$

3.2 Local Value

Unfortunately, the value function in Eq. (1) is not locally plannable by the agents as they have no access to the global swarm state s. From a control perspective, we therefore require an alternative notion of optimality that is solely based on the local observations and hence computable by the agents (see Section 4). As it seems natural for a planning agent to choose its actions by considering all possible swarm states that are in agreement with its local observation of the system, we define the *local value function* of agent n as

$$V_t^{(n)}(o \mid \pi) \coloneqq \mathbb{E}_{P_t(s \mid o^{(n)} = o, \pi)}\left[V^{(n)}(s \mid \pi)\right].$$

Note that, in case of a single agent, this is equivalent to the belief-state value in a partially observable Markov decision process [18, 16]. In our next proposition, we highlight two key properties of this quantitity: 1) It is not only locally plannable but also reduces the multi-agent problem to a single-agent one in the sense that all local values coincide. 2) In contrast to the private value, the local value is time-dependent because the conditional probabilities $P_t(s \mid o^{(n)} = o, \pi)$, in general, depend on time. However, it converges to a stationary value asymptotically under suitable conditions.

Proposition 1. *Consider a swarMDP as defined above and the stochastic process $\{S_t\}_{t=0}^{\infty}$ of the swarm state induced by the local policy π. If the initial state distribution $P_0(s)$ is invariant under permutation*[2] *of agents, all local value functions become identical,*

$$V_t^{(m)}(o \mid \pi) = V_t^{(n)}(o \mid \pi) \quad \forall m, n.$$

In this case, we may drop the agent index and denote the common local value function as $V_t(o \mid \pi)$. If, furthermore, it holds that $S_t \xrightarrow{a.s.} S$ (almost surely) for some S and the value function is continuous and bounded above, then the local value function $V_t(o \mid \pi)$ will converge to a limit,

$$V_t(o \mid \pi) \to V(o \mid \pi),$$

where $V(o \mid \pi) = \mathbb{E}_{P(s \mid o^{(n)} = o, \pi)}\left[V^{(n)}(s \mid \pi)\right]$.

Proof. Fix two agents, say agent 1 and 2. For these agents, define a permutation operation $\sigma : \mathcal{S}^N \to \mathcal{S}^N$ as

$$\sigma(s) \coloneqq (s^{(2)}, s^{(1)}, s^{(3)}, \ldots, s^{(N)}),$$

where $s = (s^{(1)}, s^{(2)}, s^{(3)}, \ldots, s^{(N)})$. Due to the homogeneity of the system, i.e. since $R(\xi^{(1)}(s)) = R(\xi^{(2)}(\sigma(s)))$ and $P(\tilde{s} \mid s, \pi) = P(\sigma(\tilde{s}) \mid \sigma(s), \pi)$, it immediately follows that $V^{(1)}(s \mid \pi) = V^{(2)}(\sigma(s) \mid \pi)\ \forall s$. This essentially means: the value that agent 1 gets at swarm state s is the same as the one that agent 2 would have got, had it been interchanged with agent 1, i.e. at swarm state $\sigma(s)$. Note that this is effectively the same as renaming the agents. The homogeneity of the system ensures that the symmetry of the initial state distribution $P_0(s)$ is maintained at all subsequent points in time, i.e. $P_t(s \mid \pi) = P_t(\sigma(s) \mid \pi)\ \forall s, t$. In particular, it holds that $P_t(s \mid o^{(1)} = o, \pi) = P_t(\sigma(s) \mid o^{(2)} = o, \pi)\ \forall s, t$ and, accordingly,

$$\begin{aligned} &V_t^{(1)}(o \mid \pi) - V_t^{(2)}(o \mid \pi) \\ &= \mathbb{E}_{P_t(s \mid o^{(1)} = o, \pi)}\left[V^{(1)}(s \mid \pi)\right] - \mathbb{E}_{P_t(s \mid o^{(2)} = o, \pi)}\left[V^{(2)}(s \mid \pi)\right] \\ &= \sum_s \Big(P_t(s \mid o^{(1)} = o, \pi) V^{(1)}(s \mid \pi)\ \ldots \\ &\qquad \ldots\ -\ P_t(\sigma(s) \mid o^{(2)} = o, \pi) V^{(2)}(\sigma(s) \mid \pi)\Big) = 0, \end{aligned}$$

[2] As we assume the agents in our model to be identical and hence interchangeable, it naturally follows to consider only the class of invariant distributions.

which shows that the local value functions are indeed identical for all agents. Now, treating the value as a random variable and using the fact that it is a continuous function, it follows that $V^{(n)}(S_t \mid \pi) \xrightarrow{a.s.} V^{(n)}(S \mid \pi)$ since $S_t \xrightarrow{a.s.} S$. As we assume the function to be finite, i.e. $|V^{(n)}(S_t \mid \pi)| < V^* \text{ for some } V^* \in \mathbb{R}$, it holds by conditional dominated convergence theorem [5] that $\mathbb{E}\left[V^{(n)}(S_t \mid \pi) \mid o^{(n)} = o, \pi\right] \to \mathbb{E}\left[V^{(n)}(S \mid \pi) \mid o^{(n)} = o, \pi\right]$, i.e. $V_t(o \mid \pi) \to V(o \mid \pi)$.

□

3.3 Global Value

The (asymptotic) local value derived in the last section provides a means to assess the quality of a certain policy from an agent perspective and can hence be used to realize the policy update step in Algorithm 1 (see details in Section 4). For the value estimation step, however, we require a measure of global system performance which allows us to compare the expert behavior with the learned one and to update the reward model accordingly. Therefore, we extend the idea presented in [22] and assign a single *global value* to agent n, representing its average performance under π

$$V_{|\pi}^{(n)} \coloneqq \mathbb{E}_{P_0(s)}\left[V^{(n)}(s \mid \pi)\right].$$

From the symmetry, i.e. since $P_0(s) = P_0(\sigma(s))$, it immediately follows that this value is independent of the specific agent under consideration, i.e.

$$V_{|\pi}^{(m)} = V_{|\pi}^{(n)} \quad \forall m, n.$$

Accordingly, we can construct an unbiased estimator for this quantity from any local trajectory of the system whose initial swarm state s_0 is drawn from $P_0(s)$, i.e.

$$\hat{V}_{|\pi}^{(n)} = \sum_{t=0}^{\infty} \gamma^t R(\xi^{(n)}(s_t)) = \sum_{t=0}^{\infty} \gamma^t r_t^{(n)}. \tag{2}$$

As these local estimates are identically distributed for all agents, we can increase the accuracy of the resulting estimate by considering the information provided by the whole swarm,

$$\hat{V}_{|\pi} = \frac{1}{N}\sum_{n=1}^{N} \hat{V}_{|\pi}^{(n)} = \frac{1}{N}\sum_{n=1}^{N}\sum_{t=0}^{\infty} \gamma^t r_t^{(n)}. \tag{3}$$

Note, however, that this is not the same as averaging N independent Monte Carlo draws from the process as the agents are correlated due to their local coupling. Nevertheless, the resulting estimator remains unbiased and, therefore, we should exploit the additional information provided by the other agents. Also, we shall see in Section 5 that, depending on the coupling strength of the system, the correlation between the agents can drop rapidly with their topological distance.

Even though the value in Eq. (3) is computed from local agent trajectories, it can be interpreted as a measure for the overall performance of the system (as dictated by the underlying reward function) in the form of a single numerical value - just as required for the IRL procedure. An alternative interpretation of the formula which shows this correspondence more explicitly is the following:

A swarm-level interpretation

Instead of starting with the private value of an individual agent as done in Eq. (1), we could define a swarm-level value function as

$$\bar{V}(s \mid \pi) \coloneqq \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k \bar{R}(s_{t+k}) \mid s_t = s, \pi\right],$$

where $\bar{R} : \mathcal{S}^N \to \mathbb{R}$ is a swarm-level reward function assessing global behavior. In an identical agent architecture, such as a swarm, it makes sense to assume that the global reward is composed of local rewards in a symmetric way, i.e. $\bar{R}(s) = \frac{1}{N}\sum_{n=1}^{N} R(\xi^{(n)}(s))$. The above swarm-level value function can then be expressed in terms of the private value functions,

$$\bar{V}(s \mid \pi) = \frac{1}{N}\sum_{n=1}^{N} V^{(n)}(s \mid \pi),$$

which essentially provides an alternative view on the value estimation in Eq. (3).

4 Optimal Control in Swarm Systems

Having defined a suitable performance measure for the system, we now move on and focus on the resulting control problem, that is, we target the problem of finding a (near-)optimal local policy for a given reward function. Besides its algorithmic requirement as part of the overall IRL procedure, we also have a primary interest in the system policy as a second key quantity - next to the reward function - to understand the observed swarm behavior. In system identification and distributed system design, for example, this knowledge is important as the policy tells us how the global dynamics of a distributed systems are realized in terms of local interactions at the agent-level. As already pointed out in Section 3, we herein restrict our search to the class of reactive policies and, based on the model properties derived in the last sections, present a learning scheme that is particularly tailored to the swarm setting.

Similar to the local value, we start by defining a local *local Q-function* for each agent, i.e.

$$Q_t^{(n)}(o, a \mid \pi) \coloneqq \mathbb{E}_{P_t(s \mid o^{(n)} = o, \pi)}\left[Q^{(n)}(s, a \mid \pi)\right],$$

which can be used to assess the quality of a particular action by quantifying the corresponding expected return of the respective agent. Following a similar argument as before, one can show that these Q-functions are again identical for all agents and, moreover, that they converge to the following value,

$$Q(o, a \mid \pi) = \mathbb{E}_{P(s \mid o^{(n)} = o, \pi)}\left[Q^{(n)}(s, a \mid \pi)\right], \tag{4}$$

which can be understood as the state-action value of a generic agent that is coupled to a stationary field generated by (and executing) policy π. Accordingly, we define an agent's optimal response to this field as

$$\pi_{\mathrm{R}}(o \mid \pi) \coloneqq \arg\max_a Q(o, a \mid \pi)$$

and observe that any optimal system policy π^*, by definition, is a fixed-point of this equation, i.e.

$$\pi_{\mathrm{R}}(o \mid \pi^*) = \pi^*(o).$$

This brings us to the idea of framing our control problem as a fixed-point problem whose solution is reached as soon as the upper condition is fulfilled. More precisely, we propose the following iterative

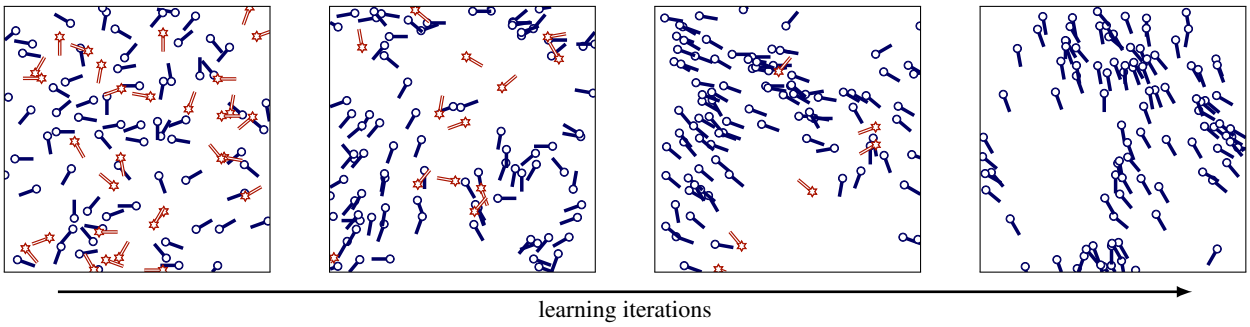

Figure 2: Illustration of the proposed reinforcement learning scheme applied to the Vicsek model (Section 5.1). The agents are divided into a greedy set ( ) and an exploration set ( ) in order to facilitate the exploration at desynchronized swarm states. The size of the exploration set is reduced over time to gradually transfer the system into a stationary behavior.

learning algorithm: At each iteration, we update the system's policy according to the best response to the current estimate of the stationary Q-function. The policy is then used to generate a new swarm behavior which, in turn, defines a corresponding Q-function. These two steps are repeated until we eventually reach a steady state. In practice, we might not be able to wait until the system reaches its stationary behavior, which is why we blend both steps into a single one (i.e. we run the system in parallel to the Q-update). Essentially, such a procedure is analogous to a model-free variant of policy iteration [14], with the important difference to the single-agent case that the environment of each agent becomes non-stationary due to the repeated altering of the swarm behavior. Therefore, we require the Q-update procedure to obey some sort of "forgetting-property" in order to limit the influence of experience acquired at early stages before reaching the optimal behavior at the fix-point.[3]

While at the fixed-point, the agents finally reach a consistent, locally optimum behavior in the form of a symmetric Nash equilibrium, we still have to ensure that the system eventually converges towards such a solution. Now the problem is that, while no agent initially performs optimally, there is no reference for any of the agents to learn from. To mitigate this problem, we propose the following heterogeneous learning scheme:

A Heterogeneous Cooling Approach

To facilitate the state exploration, we separate the agents into two disjoint sets: a greedy set and an exploration set. While the agents in the exploration set randomly explore the effect of different actions, the agents in the greedy set provide a reference behavior in the form of the optimal response to the current Q-function shared among all agents. Over time, more and more exploring agents are assigned to the greedy set so that the system is gradually transferred into a stationary behavior and thereby smoothly guided towards a fixed-point policy (see Figure 2). The heterogeneity of the learning scheme ensures that also swarm states which are locally desynchronized (e.g. when a new agent joins a group of other agents) are well-explored so that the agents can learn adequate responses to out-of-equilibrium situations. This phenomenon is best illustrated by the agent constellation in Figure 2 (second from right), a situation which is highly unlikely under a homogeneous learning scheme as it requires a series of subsequent exploration steps by only a few agents while all their neighbors need to behave consistently optimally at the same time.

[3] One such update procedure is realized, for instance, by the Q-learning algorithm [33], which brings this property naturally in the form of its learning rate.

The final algorithm is summarized below, together with a pictorial description of the main steps. Please note that, while such a cooling approach makes sense intuitively and has been successfully applied in other contexts [29], we cannot provide any proof of convergence at this stage. Nevertheless, the simulation results in the following sections demonstrate that we indeed achieve a system performance comparable to the that of the expert policy.

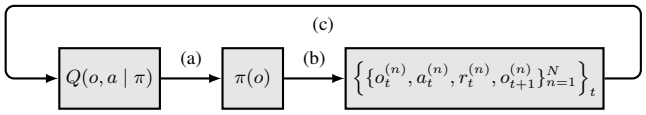

Figure 3: Pictorial description of the swarm learning scheme. (a) The next system policy is estimated from the current guess of the stationary Q-function. (b) Heterogeneous one-step transition of the system. (c) The estimate of the Q-function is updated, with past samples having less influence ("forgetting-property").

Algorithm 2: HETEROGENEOUS Q-LEARNING

Input: swarMDP without policy π

1: Initialize shared Q-function, learning rate and fraction of exploring agents (i.e. the temperature)
for $i = 0, 1, 2, ...$
2: Separate the swarm into exploring and greedy agents according to the current temperature
3: Based on the current swarm state and Q-function, select actions for all agents
4: Iterate the system and collect rewards
5: Update the shared Q-function based on the new experience
6: Decrease the learning-rate and the temperature
end for

5 Simulation Results

Due to the lack of real data, we present simulation results for our framework on two different models. In both cases, we used the max-margin approach presented in [1] as the underlying IRL framework. The required value estimates were obtained from a single system trajectory as indicated by Eq. (3). Due to its convenient algorithmic form which naturally fulfills our requirements on the policy update, we used a simple variant of Q-learning in which the agents' common

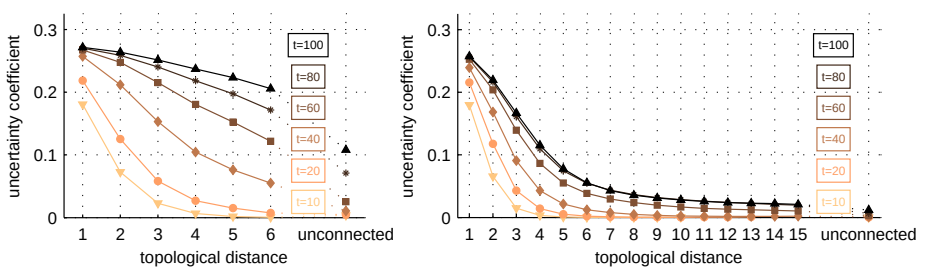

(a) Uncertainty coefficient for interaction radii $\rho = 0.125$ (left) and $\rho = 0.05$ (right). The values are based on a kernel density estimate of the joint distribution of two agents' headings.

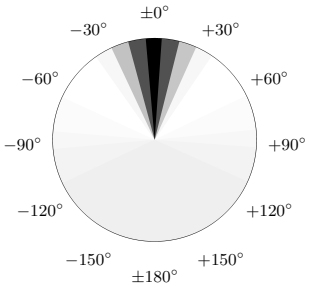

(b) Mean reward function.

Figure 4: Results for the Vicsek model. All values are estimated from 10000 Monte Carlo runs. (a) Uncertainty coefficient as a function of the agents' topological distance. (b) Mean reward as a function of local misalignment learned by the algorithm.

Q-function (Eq. (4)), stored as a table, was synchronously updated by all agents according to the learning scheme proposed in Algorithm 2. The initial number of exploring agents was set to 50% of the population size and the initial learning rate was set to 1. Both quantities were controlled by a quadratic decay which ensured that, at the end of the learning period, i.e. after 200 iterations, the learning rate reached zero and there were no exploring agents left. Please note that these parameters were by no means optimized. However, during our simulations we could observe that the final results were barely affected by the choice of these parameters for a wide range.

5.1 The Vicsek Model

In our first experiment, we tested a modified version of the so-called *Vicsek* or *self-propelled particles* model [30]. It consists of a fixed number of particles or agents living in the unit square $[0, 1] \times [0, 1]$ with periodic boundary conditions. All agents share a common absolute velocity v but every agent has its individual orientation $\theta_t^{(n)}$ in the plane. At each iteration, the agents' orientations get synchronously updated to the average orientation of their neighbors (including themselves), while the time-varying neighborhood structure of the system is determined by a predefined interaction radius ρ. After this update, the resulting orientations are distorted by random perturbations, in the following denoted as $\Delta\theta_t^{(n)}$:

$$\begin{aligned} \theta_{t+1}^{(n)} &= \langle \theta_t^{(n)} \rangle_\rho + \Delta\theta_t^{(n)}, \\ x_{t+1}^{(n)} &= x_t^{(n)} + v_t^{(n)}. \end{aligned} \tag{5}$$

Here, $\langle \theta_t^{(n)} \rangle_\rho$ denotes the mean orientation of all particles within the ρ-neighborhood of particle n, $x_t^{(n)}$ describes its position in the plane and $v_t^{(n)} = v \cdot [\cos\theta_t^{(n)}, \sin\theta_t^{(n)}]$ is its current velocity vector. Accordingly, we define an agent's local state as $s_t^{(n)} := (x_t^{(n)}, \theta_t^{(n)})$.

As the original Vicsek dynamics can generally lead to arbitrary changes in direction, the exact behavior of the system can not be described by a finite set of actions and would require a continuous-action MDP model. Since we are, however, only interested in replicating the basic dynamics of the system, we discretized the space of possible direction changes to $[-60°, -50°, \ldots, 60°]$, resulting in a total of 13 actions. The expert policy is defined such that the resulting swarm system most closely resembles the continuous dynamics (i.e. we always select the action with smallest angular distance to the ideal change in direction). As a simple observation model, we used the relative angles between the agents' orientations and the average orientation of their neighbors, i.e. $o_t^{(n)} = \xi^{(n)}(s_t) = \langle \theta_t^{(n)} \rangle_\rho - \theta_t^{(n)}$, giving the agents the ability to monitor their local misalignment. As the space of possible observations grows exponentially with the number of agents, we furthermore coarse-grained the observation space using a finite number of basis functions as suggested in [22]. To this end, we divided the interval $[0, 2\pi)$ into 36 center positions and placed a copy of a rectangular basis function prototype of $10°$ width at each position to indicate whether an observation falls into a certain range. For the simulations, we used a system size of $N = 200$ agents, an interaction radius of $\rho = 0.1$ (if not stated otherwise), an absolute velocity of $v = 0.1$, a discount factor of $\gamma = 0.9$ and a zero-mean Gaussian noise model for the perturbations $\Delta\theta_t^{(n)}$ with a standard deviation of $10°$.

Local Coupling & Redundancy

In Section 3.3 we claimed that, due to the local coupling structure of a swarm system, the correlation between two agents will drop rapidly with their topological distance so that we can gain information by monitoring the whole swarm instead of looking at just a single agent. In this section, we shall analyze the resulting information gain for the Vicsek model by measuring the redundancy in the system as a function of the topological distance between the agents for different coupling strengths.

As a measure of (in)dependence, we used the so-called *uncertainty coefficient* [24], a normalized version of the mutual information shared between two agents, which tells us the fraction of bits we can predict about one agent's orientation by observing that of the other. As opposed to linear correlation, this measure is able to capture non-linear dependencies and is hence more meaningful in the context of the Vicsek model whose angular state dynamics are inherently non-linear. Figure 4(a) depicts the result of our analysis which nicely reveals the flow of information throughout the system over time. As expected, the mutual dependence between the agents strongly depends on the strength of their coupling which is determined by 1) their topological distance and 2) the number of connecting links (seen from the fact that, for a fixed distance, the dependence grows with the interaction radius). We can also see that, for increasing interaction radii, the dependence grows even for agents which are temporarily not connected through the system due to the increasing chances of having been connected at some earlier stage.

Learning Results

An inherent problem with any IRL approach is the assessment of the extracted reward function as there is typically no ground truth

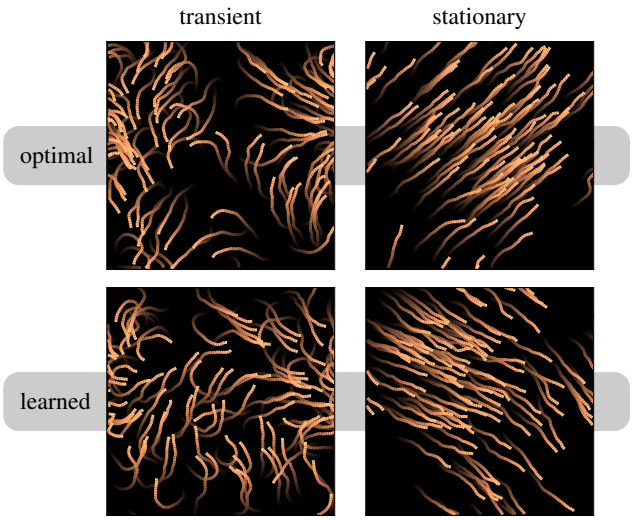

Figure 5: Illustrative trajectories of the Vicsek model generated under the optimal and a learned policy. A color-coding scheme is used to indicate the temporal progress.

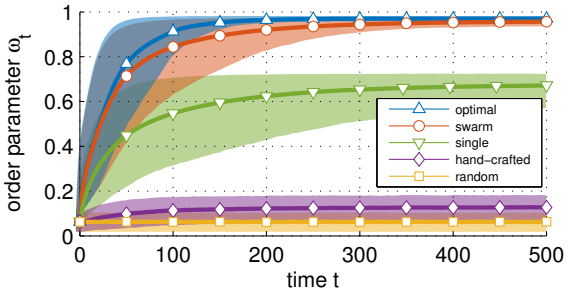

Figure 6: Temporal evolution of the order parameter ω_t of the Vicsek model. From top to bottom, the curves show the results for the expert policy, for the proposed swarm IRL approach (i.e. the trajectories of all agents are evaluated), the result we get when values are estimated from just one single agent, for a hand-crafted reward function, and for random policies. For the optimal policy, we show the empirical mean and the corresponding 10% and 90% quantiles, based on 10000 Monte Carlo runs. For the learned policies - since the outcome of the learning process is random - we instead show the average over 100 conditional quantiles, each based on 100 Monte Carlo runs.

to compare with. The simplest way to check the tenability of the result is of course by subjective inspection: since the reward can be regarded as a concise description of the task being performed, the reward function should be able to explain the observed system behavior reasonably well. As we can see from Figure 4(b), this is indeed the case for the found result. Although the true reward of the model is unknown, we can see from the system equations in (5) that particles tend to align over time. Clearly, such dynamics can be induced by giving higher rewards for synchronous behavior and low or negative reward for misalignment. Additionally, we can inspect the induced behavior itself, as shown in Figure 5. We can clearly see that the algorithm is able to capture the basic behaviorial patterns of the swarm, both during the transient phase and at stationarity. Please note that the absolute direction of travel is not important here as the model considers only relative angles between the agents.

Finally, we compare the results in terms of the *order parameter* [30], which provides a measure for the total alignment of the swarm,

$$\omega_t := \frac{1}{Nv}\left|\sum_{n=1}^{N} v_t^{(n)}\right| \in [0,1],$$

with values close to 1 indicating strong synchronization of the system. In Figure 6, we show the temporal evolution of this quantity under the true policy, pitted against the learned ones. From the result, we can see a considerable gain in performance for the proposed estimation scheme (Eq. (3)) when compared to the single-agent approach (Eq. (2)), which obviously stems from the extra information provided by the other agents. As additional reference, the figure also shows the result for using a hand-crafted reward model, where positive reward is provided only by the basis function containing the synchronized state. As we can see, the learned reward model significantly outperforms the ad-hoc solution.

5.2 The Ising Model

In a second experiment, we applied our framework to the well-known *Ising model* [13] which, in our case, consists of a finite grid of atoms (or agents) of size 100×100. Each agent has its individual spin $q_t^{(n)} \in \{+1, -1\}$ which, together with its position on the grid, forms its local state, i.e. $s_t^{(n)} := (x^{(n)}, y^{(n)}, q_t^{(n)})$. For our simulations, we considered a static 5×5-neighborhood system meaning that each agent interacts only with its 24 closest neighbors, i.e. agents with a maximum Chebyshev distance[4] of 2. Based on this neighborhood relationship, we define the total energy of the system as

$$E_t := \sum_{n=1}^{N} \sum_{m \in \mathcal{N}_n} \mathbb{1}(q_t^{(n)} \neq q_t^{(m)}) = \sum_{n=1}^{N} E_t^{(n)},$$

where $\mathbb{1}(\cdot)$ denotes the indicator function and $\mathcal{N}_n$ and $E_t^{(n)}$ are the neighborhood of agent n and its local contribution to the total energy, respectively. Just as the order parameter for the Vicsek model, the energy serves as a measure for the total alignment of the system, with zero energy indicating complete synchronization. For our experiment, we considered two possible actions that are available to the agents, i.e. *keep the current spin* and *flip the spin*. The local transition model was chosen such that, depending on the chosen action, an agent ended up in the desired next state with probability 1. As before, we gave the agents the ability to monitor only their local misalignments which, in this case, is given by their local energy contributions, i.e. $o_t^{(n)} = \xi^{(n)}(s_t) = E_t^{(n)}$.

Now, a meaningful goal for the system is to reach a global state configuration of minimum energy and we are interested in finding a *local* policy which can achieve this goal. Instead of learning such behavior using a hand-crafted reward model, we prefer to extract the reward function from expert demonstrations using our IRL framework. As expert behavior, we choose the local policy which always decides for the action that would result in the lower local energy level if the neighboring agents did not change their spins, corresponding to a local majority vote. Essentially, this is equivalent to a synchronous version of the *iterated conditional modes* algorithm [4] which, in the sequential case, is guaranteed to translate the system to a state of locally minimum energy.

Figure 7 and 8 show, respectively, the mean reward function found by the algorithm and the evolution of the global energy under the

[4] The Chebyshev distance between two points $(x, y), (x', y') \in \mathbb{R}^2$ is defined as $D((x, y), (x', y')) := \max(|(x - x')|, |(y - y')|)$.

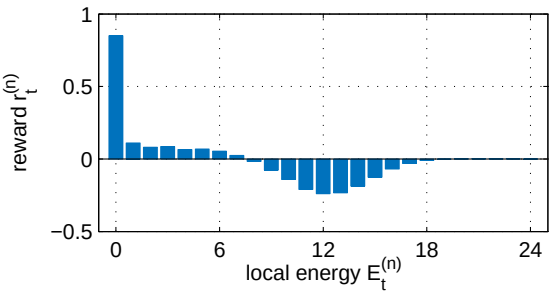

Figure 7: Mean reward as a function of local energy learned from the Ising model in 10000 Monte Carlo runs.

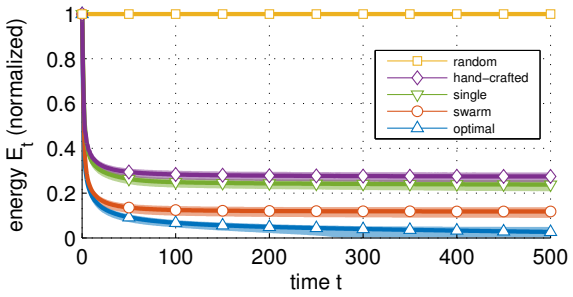

Figure 8: Temporal evolution of the total energy E_t in the Ising model. The graphs are analogous to those in Figure 6.

learned policies. As for the Vicsek model, the extracted reward function is reasonable in terms of the observed system behavior[5] and we can observe the same qualitative improvement in terms of the system energy as for the order parameter, both when compared to the single-agent estimation scheme and to the hand-crafted reward model.

6 Discussion

Our objective in this paper has been to apply the concept of IRL to homogeneous multi-agent systems, called swarms, in order to learn a local reward function from observed global system dynamics that could explain the emergent behavior of the system. We showed that, by exploiting the inherent homogeneity of the swarm, both value estimation and policy update required for the IRL procedure can be efficiently performed based on local experience at the agent-level which, in the end, allowed us to synthesize the system's global behavior in terms of local interactions. We substantiated our claims by applying the developed framework to the Vicsek model and the Ising model, where we could show a performance close to the observed expert behavior in terms of system synchronization.

Nevertheless, there remain some open questions. For the purpose of value estimation, we have tacitly assumed that the agents merely follow the system's instructions encoded in the form of local policies provided to them. For an arbitrary system, however, we can not ascertain whether the behavior was planned on a global or a local level, which brings us to the following questions: Which value function is actually optimized by the system? How clear are the agents about the global behavior of the swarm? Is it always possible to reconstruct an observed behavior, perhaps planned on the swarm level, by optimizing local values? If not, how well can we approximate it? In both our experiments, we were in fact able to do so but it remains unclear when local planning will not suffice to reproduce the global behavior.

In an attempt to understand the above mentioned problem, we propose the following characterization of the reward function that would make a local policy optimal in a swarm. To this end, we enumerate the swarm states and observations by $\mathcal{S}^N = \{s_i\}_{i=1}^K$ and $\mathcal{O} = \{o_i\}_{i=1}^L$, respectively. Furthermore, we fix an agent n and define matrices $\mathbf{P}_o, \{\mathbf{P}_a\}_{a=1}^{|\mathcal{A}|}$, where $[\mathbf{P}_o]_{ij} = P(s_j \mid o^{(n)} = o_i, \pi)$ and $[\mathbf{P}_a]_{ij} = P^{(n)}(s_j \mid s_i, a, \pi)$. Finally, we express the reward as a vector, i.e. $\mathbf{R} = \big(R(\xi^{(n)}(s_1)), \dots, R(\xi^{(n)}(s_K))\big)^\mathsf{T}$.

Proposition 2. *Consider a swarm $(N, \mathbb{A}, \xi)$ of agents $\mathbb{A} = (\mathcal{S}, \mathcal{O}, \mathcal{A}, T, R, \pi)$ and a discount factor $\gamma \in [0, 1)$. Then a policy $\pi : \mathcal{O} \to \mathcal{A}$ given by $\pi(o) := a_1$ is optimal[6] with respect to $V(o \mid \pi)$ if and only if the reward R satisfies*

$$\mathbf{P}_o(\mathbf{P}_{a_1} - \mathbf{P}_a)(\mathbf{I} - \gamma\mathbf{P}_{a_1})\mathbf{R} \geq 0 \quad \forall a \in \mathcal{A}, \tag{6}$$

where $\geq$ is meant component-wise.

Proof. Expressing the private value in Eq. (1) using vector notation, we have

$$\mathbf{V}_{s|\pi} = (\mathbf{I} - \gamma\mathbf{P}_{a_1})^{-1}\mathbf{R},$$

where $\mathbf{V}_{s|\pi} = \big(V^{(n)}(s_1 \mid \pi), \dots, V^{(n)}(s_K \mid \pi)\big)^\mathsf{T}$. According to Proposition 1, the corresponding limiting local value function is given by

$$V(o \mid \pi) = \sum_{i=1}^K P(s_i \mid o^{(n)} = o, \pi) V^{(n)}(s_i \mid \pi).$$

Rewritten in vector notation, we obtain

$$\mathbf{V}_{o|\pi} = \mathbf{P}_o \mathbf{V}_{s|\pi}, \tag{7}$$

where $\mathbf{V}_{o|\pi} = \big(V(o_1 \mid \pi), \dots, V(o_L \mid \pi)\big)^\mathsf{T}$. Now, $\pi(o) = a_1$ is optimal if and only if for all $a \in \mathcal{A}, o \in \mathcal{O}$

$$\begin{aligned}
&Q^{(n)}(o, a_1 \mid \pi) \geq Q^{(n)}(o, a \mid \pi) \\
&\Leftrightarrow \sum_{i=1}^K P(s_i \mid o^{(n)} = o, \pi) \sum_{j=1}^K P^{(n)}(s_j \mid s_i, a_1, \pi) V^{(n)}(s_j \mid \pi) \\
&\quad \geq \sum_{i=1}^K P(s_i \mid o^{(n)} = o, \pi) \sum_{j=1}^K P^{(n)}(s_j \mid s_i, a, \pi) V^{(n)}(s_j \mid \pi) \\
&\Leftrightarrow \mathbf{P}_o(\mathbf{P}_{a_1} - \mathbf{P}_a)\mathbf{V}_{s|\pi} \geq 0 \\
&\Leftrightarrow \mathbf{P}_o(\mathbf{P}_{a_1} - \mathbf{P}_a)(\mathbf{I} - \gamma\mathbf{P}_{a_1})^{-1}\mathbf{R} \geq 0,
\end{aligned}$$

□

Remark. *Following a similar derivation as in [22], we obtain the characterization set with respect to $\mathbf{V}_{s|\pi}$ as*

$$(\mathbf{P}_{a_1} - \mathbf{P}_a)(\mathbf{I} - \gamma\mathbf{P}_{a_1})\mathbf{R} \geq 0. \tag{8}$$

Notice that, as Eq. (8) implies Eq. (6), an $\mathbf{R}$ that makes $\pi(o)$ optimal for $\mathbf{V}_{s|\pi}$, also makes it optimal for $\mathbf{V}_{o|\pi}$. Therefore, denoting by $\mathcal{R}_L$ and $\mathcal{R}_G$ the solution sets corresponding to the local and global values $\mathbf{V}_{o|\pi}$ and $\mathbf{V}_{s|\pi}$, respectively, we conclude

$$\mathcal{R}_G \subseteq \mathcal{R}_L,$$

with equality in the trivial case where observation o is sufficient to determine swarm state s. It is therefore immediate that, as long as there is uncertainty about the swarm state, local planning can only guarantee globally optimal behavior in an average sense as pronounced by $\mathbf{P}_o$ (see Eq. (7)).

[5] Please note that assigning a neutral reward to states of high local energy is reasonable, since strong local misalignment indicates a high synchronization of opposite spin for the neighbors.

[6] We can always ensure that $\pi(o) = a_1$ is fulfilled by renaming actions (see [22]).

REFERENCES

[1] P. Abbeel and A. Y. Ng, 'Apprenticeship learning via inverse reinforcement learning', in *Proc. 21st International Conference on Machine Learning (ICML)*, p. 1. ACM, (2004).

[2] A. Agogino and K. Tumer, 'Reinforcement learning in large multi-agent systems', in *Proc. AAMAS 2005 Workshop on Coordination of Large Scale Multi-Agent Systems*, (2005).

[3] K. C. Berridge and M. L. Kringelbach, 'Affective neuroscience of pleasure: reward in humans and animals', *Psychopharmacology*, **199**(3), 457–480, (2008).

[4] J. Besag, 'On the statistical analysis of dirty pictures', *Journal of the Royal Statistical Society. Series B (Methodological)*, 259–302, (1986).

[5] P. Billingsley, *Convergence of probability measures*, John Wiley & Sons, 2013.

[6] J. Buhl, D. Sumpter, I. D. Couzin, J. J. Hale, E. Despland, E. Miller, and S. J. Simpson, 'From disorder to order in marching locusts', *Science*, **312**(5778), 1402–1406, (2006).

[7] L. Busoniu, R. Babuska, and B. De Schutter, 'A comprehensive survey of multiagent reinforcement learning', *Proc. IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, **38**(2), 156–172, (2008).

[8] I. D. Couzin, 'Collective cognition in animal groups', *Trends in cognitive sciences*, **13**(1), 36–43, (2009).

[9] L. Dufton and K. Larson, 'Multiagent policy teaching', in *Proc. International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, (2009).

[10] R. A. Freitas, 'Current status of nanomedicine and medical nanorobotics', *Journal of Computational and Theoretical Nanoscience*, **2**(1), 1–25, (2005).

[11] S. C. Goldstein, J. D. Campbell, and T. C. Mowry, 'Programmable matter', *Computer*, **38**(6), 99–101, (2005).

[12] M. Grzes and D. Kudenko, 'Learning shaping rewards in model-based reinforcement learning', in *Proc. AAMAS 2009 Workshop on Adaptive Learning Agents*, volume 115, (2009).

[13] E. Ising, 'Beitrag zur Theorie des Ferromagnetismus', *Zeitschrift für Physik*, **31**(1), 253–258, (1925).

[14] M. G. Lagoudakis and R. Parr, 'Least-squares policy iteration', *The Journal of Machine Learning Research*, **4**, 1107–1149, (2003).

[15] V. Lesser, Charles L. Ortiz J., and M. Tambe, *Distributed sensor networks: a multiagent perspective*, volume 9, Springer Science & Business Media, 2012.

[16] F. S. Melo, 'Exploiting locality of interactions using a policy-gradient approach in multiagent learning', in *Proc. 18th European Conference on Artificial Intelligence*, p. 157, (2008).

[17] F. S. Melo, M. Lopes, and R. Ferreira, 'Analysis of inverse reinforcement learning with perturbed demonstrations', in *Proc. 19th European Conference on Artificial Intelligence*, pp. 349–354, (2010).

[18] N. Meuleau, L. Peshkin, K.-E. Kim, and L. P. Kaelbling, 'Learning finite-state controllers for partially observable environments', in *Proc. 15th Conference on Uncertainty in Artificial Intelligence*, pp. 427–436, (1999).

[19] B. Michini and J. P. How, 'Bayesian nonparametric inverse reinforcement learning', in *Machine Learning and Knowledge Discovery in Databases*, 148–163, Springer, (2012).

[20] S. Natarajan, G. Kunapuli, K. Judah, P. Tadepalli, K. Kersting, and J. Shavlik, 'Multi-agent inverse reinforcement learning', in *Proc. 9th International Conference on Machine Learning and Applications (ICMLA)*, pp. 395–400. IEEE, (2010).

[21] A. Y. Ng, D. Harada, and S. Russell, 'Policy invariance under reward transformations: theory and application to reward shaping', in *Proc. 16th International Conference on Machine Learning (ICML)*, volume 99, pp. 278–287, (1999).

[22] A. Y. Ng and S. Russell, 'Algorithms for inverse reinforcement learning', in *Proc. 17th International Conference on Machine Learning (ICML)*, pp. 663–670, (2000).

[23] E. Omel'chenko, Y. L. Maistrenko, and P. A. Tass, 'Chimera states: the natural link between coherence and incoherence', *Physical review letters*, **100**(4), 044105, (2008).

[24] W. H. Press, *Numerical recipes 3rd edition: The art of scientific computing*, Cambridge University Press, 2007.

[25] T. S. Reddy, V. Gopikrishna, G. Zaruba, and M. Huber, 'Inverse reinforcement learning for decentralized non-cooperative multiagent systems', in *Proc. IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 1930–1935. IEEE, (2012).

[26] C. A. Rothkopf and C. Dimitrakakis, 'Preference elicitation and inverse reinforcement learning', in *Machine Learning and Knowledge Discovery in Databases*, 34–48, Springer, (2011).

[27] Y. Shoham, R. Powers, and T. Grenager, 'Multi-agent reinforcement learning: a critical survey. Technical report', (2003).

[28] R. S. Sutton and A. G. Barto, *Reinforcement learning: an introduction*, volume 1, MIT press Cambridge, 1998.

[29] P. J. van Laarhoven and E. H. Aarts, *Simulated annealing: theory and applications*, volume 37, Springer Science & Business Media, 1987.

[30] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, 'Novel type of phase transition in a system of self-driven particles', *Physical review letters*, **75**(6), 1226, (1995).

[31] T. Vicsek and A. Zafeiris, 'Collective motion', *Physics Reports*, **517**(3), 71–140, (2012).

[32] T. Voice, R. Stranders, A. Rogers, and N. R. Jennings, 'A hybrid continuous max-sum algorithm for decentralised coordination', in *Proc. 19th European Conference on Artificial Intelligence*, pp. 61–66, (2010).

[33] C. Watkins and P. Dayan, 'Q-learning', *Machine learning*, **8**(3-4), 279–292, (1992).

[34] G. M. Whitesides and B. Grzybowski, 'Self-assembly at all scales', *Science*, **295**(5564), 2418–2421, (2002).

[35] S. Zhifei and E. M. Joo, 'A survey of inverse reinforcement learning techniques', *International Journal of Intelligent Computing and Cybernetics*, **5**(3), 293–311, (2012).

[36] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, 'Maximum entropy inverse reinforcement learning', in *Proc. 23rd Conference on Artificial Intelligence (AAAI)*, pp. 1433–1438, (2008).